

Global Logic and Local Search: Dual-Stream Multimodal In-Context Learning for Verifiable Industrial Anomaly Detection

Runzhi Deng^{1*}, Yundi Hu^{1*}, Yiming Zhong^{1*}, Zhao Wang^{2†}, Xixi Liu²,
Hongsong Wang³, Caifeng Shan¹, and Fang Zhao^{1†}

¹ Nanjing University, Nanjing, Jiangsu, China

² China Mobile Zijin Innovation Institute, Nanjing, Jiangsu, China

³ Southeast University, Nanjing, Jiangsu, China

Abstract. Large Multimodal Models (LMMs) show strong few-shot generalization, but industrial anomaly detection remains difficult because defects are small, input resolution is limited, and textual standards are not always grounded in visual evidence. Recent optimization-based methods improve alignment through fine-tuning, but they often require many defective samples, which are unavailable in early deployment. We present **Global Logic and Local Search (GLLS)**, a training-free framework for reference-guided multimodal in-context verification. GLLS uses a Part-Aware Visual-Logical Atlas to organize normal references and structured specifications in the inference context. It combines a **Global & Logic Stream**, where SAM 3 extracts partially checkable visual facts, with a **Fine-Grained & Actions Stream**, where MCTS selects local evidence crops under a fixed budget. Experiments on MMAD-QA and additional anomaly detection datasets show consistent gains over matched and general-purpose baselines, while keeping the final diagnostic decision traceable to explicit visual evidence throughout the inspection trace.

Keywords: Industrial anomaly detection · Multimodal in-context learning · Multimodal large language models

1 Introduction

Industrial anomaly detection fundamentally differs from general recognition tasks because it functions as a rigorous compliance audit against explicit specifications rather than open-ended inference. Unlike open-world scenarios where models must infer normality from large-scale data, industrial inspection requires the identification of deviations from rigorous standards, such as the topology of components and geometric tolerances. This challenge is particularly acute during the initial phase of deployment, where defective samples are non-existent, normal references are scarce, and the system must rely heavily on natural language

* Equal contribution.

† Corresponding authors.




Fig. 1: Motivation and Performance Overview. The left panels visualize critical failure modes in industrial inspection, where LMMs typically suffer from **Logic Hallucination** and **Visual Bias** (top), while specialized models of smaller scale struggle with **Semantic Blindness** and **Resolution Limits** (bottom). We address these limitations by incorporating a tool-grounded *Global Logic* stream and an active *Local Search* stream to support auditable inspection. The right panels present quantitative analyses showing that the GLLS framework improves over representative general-purpose baselines under a controlled budget of computation across five metrics on the MMAD-QA benchmark.

specifications. While Large Multimodal Models (LMMs) exhibit remarkable generalization capabilities [2, 18], the direct application of these models to this task remains challenging. They inherently struggle to ground abstract textual standards into verification at the pixel level and frequently fail to resolve fine-grained defects due to inherent limitations in input resolution.

To bridge the gap between abstract standards and visual inspection, recent studies have explored various adaptation strategies. Recent optimization-based methods attempt to align LMMs via reinforcement learning, multi-agent frameworks, or extensive supervised fine-tuning to enforce consistency in reasoning [10, 14, 22, 27, 31, 43]. However, these approaches are computationally intensive and demand extensive data, necessitating numerous defective samples that are typically unavailable during the initial phases of deployment. To mitigate this data dependency, logical reasoning and retrieval-enhanced methods such as Echo [11] and Triad [28] utilize external bases of knowledge or expert-guided tokenizers [24, 26, 36]. Nevertheless, these methods typically rely on static context

injection [39] or generic region proposals, which fail to synthesize the global logic of assembly and the local sensory evidence required for auditable inspection.

In contrast, human experts manage this complexity through a cognitive process described by the *Reverse Hierarchy Theory* of vision [16]. Inspectors do not strictly rely on massive training data; instead, they consult technical specifications to form a high-level understanding of structures, and subsequently direct their attention to verify specific details via active perception [4], echoing recent MLLM evidence that explicit perception tokens enhance visual reasoning [7]. This interplay between explicit domain knowledge and active visual search allows humans to audit complex assemblies even with minimal prior exposure to visual samples. We argue that a reliable industrial inspection system can benefit from an active and knowledge-grounded workflow rather than relying on passive, end-to-end mechanisms of perception.

Toward this end, we propose a training-free framework, termed **Global Logic and Local Search (GLLS)**, which is motivated by two practical requirements in cold-start industrial inspection. First, effective inspection requires decoupling tool-grounded structural checks from probabilistic visual reasoning to ensure that logical constraints are explicitly evaluated and remain fully auditable. Second, resolving ambiguity in high-resolution images benefits from an evidence-seeking strategy that selects informative observations under a fixed budget rather than relying only on passive retrieval. GLLS treats the LMM as an auditor orchestrated by a **Part-Aware Visual-Logical Atlas (PVLA)**. This atlas leverages natural language specifications and available reference samples to construct an executable plan for inspection, providing explicit chains of evidence for the final diagnostic conclusions.

The contributions of this work are summarized as follows:

- We propose GLLS, a training-free dual-stream framework that orchestrates LMMs with auditable intermediate evidence by synergizing a **Global & Logic Stream** for tool-grounded structural checks and a **Fine-Grained & Actions Stream** for the acquisition of fine-grained sensory evidence.
- We construct the Part-Aware Visual-Logical Atlas (PVLA) to organize natural-language specifications and normal references into an explicit graph of inspection knowledge. It supports text-guided zero-shot inspection, where SAM 3 [9] grounds logical rules without visual references, and reference-guided few-shot verification when normal samples are available.
- We design a budgeted active Local Search mechanism utilizing Monte Carlo Tree Search (MCTS) [8], which selects spatially distinct evidence crops under fixed computation and context budgets without exhaustive scanning.
- The GLLS framework achieves strong performance on the MMAD-QA benchmark and demonstrates strong generalization across various industrial datasets, including MPDD, DTD, and DAGM. By strictly controlling budgets of information and computation, our framework provides a competitive, training-free alternative to optimization-based methods, exhibiting consistent efficacy across data regimes ranging from text-guided zero-shot settings to reference-guided few-shot settings.

2 Related Work

Generalized Multimodal Anomaly Detection. Current research on aligning Multimodal Models with industrial standards encompasses both optimization-based strategies and lightweight adaptations. To enforce consistent reasoning chains, reinforcement learning approaches, such as IAD-R1 [27], AnomalyR1 [10], and EMIT [14], utilize Group Relative Policy Optimization to align outputs with human logic, while LR-IAD [41] employs specialized reward functions to address class imbalance. In parallel with these RL-based methods, architectural modifications characterize frameworks such as EIAD [42] and Triad [28], which decouple dialog generation or incorporate expert tokenizers to focus on critical regions. Furthermore, extensive fine-tuning frameworks, including OmniAD [43] and AgentIAD [31], have been proposed to enhance anomaly understanding. On the spectrum of data efficiency, smaller adaptations, such as AdaptCLIP [13] and ABound [12], leverage CLIP-based architectures for boundary-driven learning, but they often lack the deep semantic reasoning capabilities required for complex diagnosis. Despite these contributions, these methods generally necessitate either extensive defective training data or costly gradient updates. Consequently, this dependency restricts their applicability during initial deployment phases, where only a few normal reference images and natural language specifications are available.

Multimodal In-Context Learning. In-Context Learning (ICL) enables models to adapt to new tasks via demonstrations without parameter updates [17]. However, standard Multimodal ICL often exhibits strong biases toward textual priors and may neglect fine-grained visual details when modalities contradict [5]. Recent research emphasizes compositional learning to improve context understanding [25] and to mitigate these hallucinations. However, generic retrieval-augmented generation and domain injection techniques [30, 33, 39] risk introducing noisy or irrelevant context, which degrades precision. We address these gaps through a reference-guided multimodal ICL framework, wherein a Part-Aware Visual-Logical Atlas injects precise executable standards into the inference context. This approach ensures that the reasoning process is grounded in auditable canonical norms rather than hallucinatory priors, thereby enabling immediate deployment without parameter updates.

3 Method

3.1 Overview of the GLLS Framework

Industrial anomaly detection demands a rigorous standard of reliability that purely data-driven foundation models often fail to meet, because a fluent description offers no guarantee that critical constraints, such as component presence or assembly order, are correctly checked. To bridge the gap between probabilistic description and auditable inspection, we propose **Global Logic and Local Search**




Fig. 2: System Overview. **Offline Construction:** The **Part-Aware Visual-Logical Atlas (PVLA)** is constructed from natural language specifications and a few normal reference images. The SAM 3 segmentation model extracts visual primitives to populate the atlas with canonical visual norms and executable logic. **Online Dual-Stream Inference:** (1) The **Global & Logic Stream** executes tool-grounded checks via the SAM 3 model to validate structural integrity. (2) In parallel, the **Fine-Grained & Actions Stream** formulates the inspection process as a budgeted Markov Decision Process (MDP), utilizing **MCTS** to actively zoom in on ambiguous regions. (3) A Vision-Language Model (VLM) performs the final **Reference-Guided In-Context Verification** by fusing structural verdicts and local evidence chains.

(GLLS), which decouples evidence acquisition from verdict reasoning. The framework uses a pre-compiled rulebook to guide the LMM as an evidence-driven auditor, rather than passively consuming images.

The workflow consists of two distinct phases. In the offline phase, we employ a hybrid pipeline to construct a **Part-Aware Visual-Logical Atlas (PVLA)** by synthesizing structured dataset priors with natural language reasoning into an executable knowledge graph. This creates a grounded memory bank where every physical component is linked to specific inspection rules. Subsequently, in the online phase, the system executes a dual-stream inference strategy. The **Global & Logic Stream** utilizes SAM 3 [9] to perform tool-grounded validation of structural logical constraints, while the **Fine-Grained & Actions Stream** uses Monte Carlo Tree Search (MCTS) as a budgeted proposal mechanism for selecting high-resolution evidence regions. By converging these streams, the final verdict is a reasoned decision grounded in both tool-grounded checks and fine-grained visual evidence.

3.2 Hybrid Atlas Construction

To ensure the framework remains scalable while maintaining strict alignment with industrial standards, we use a hybrid construction pipeline that expands sparse dataset labels into structured inspection specifications. This construction process integrates structured domain priors with the reasoning capabilities of Large Multimodal Models and the open-vocabulary segmentation power of state-of-the-art vision models.

Knowledge Expansion via GPT-5 We utilize the structured metadata provided by the MMAD benchmark as the initial logical skeleton, which provides the ground-truth set of normal component categories specific to the target domain. Since these raw labels lack the physical granularity required for detailed inspection, we employ **GPT-5** [32] to expand the seed taxonomy into candidate physical definitions based on the provided natural language ontology. Specifically, the model proposes granular sub-component attributes and potential failure modes derived from the component definitions. To ensure reliability, the generated definitions are filtered to mitigate inconsistencies. For instance, given a *Cable* component, GPT-5 infers the existence of sub-parts such as the *Outer Sheath* and *Copper Cores*, while simultaneously hypothesizing probable anomalies like *Insulation Cuts* or *Wire Deformation* based on textual context. This process helps reduce the semantic gap between abstract category names and concrete physical constraints.

Visual Grounding with SAM 3 To populate the visual nodes of the atlas with canonical reference features, we leverage the SAM 3 segmentation model [9], which provides strong text-conditioned grounding for short part phrases and produces candidate masks. We use the sub-component names generated by GPT-5 as textual prompts to guide SAM 3 in extracting regions of interest from normal reference images. To improve the reliability of these visual primitives, we filter the extracted segments using fixed confidence, IoU, and mask-stability criteria. This filtering step reduces ambiguous or low-quality proposals before they are inserted into the atlas.

Semantic Graph Formalism Formally, we define the resulting Part-Aware Visual-Logical Atlas as a heterogeneous directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ constructed to explicitly decouple canonical standards from anomalous deviations. To capture the hierarchical assembly and state duality, we formulate the vertex set $\mathcal{V}$ as the disjoint union of four semantic subsets as

$$\mathcal{V} = \{v_{cls}\} \cup \mathcal{V}_{part} \cup \mathcal{V}_{norm} \cup \mathcal{V}_{def} \quad (1)$$

At the root level, v_{cls} represents the abstract object category, which decomposes into a set of physical components $\mathcal{V}_{part} = \{v_{p_i}\}_{i=1}^N$ that mirror the topological structure of the target object. To model the ideal state, we define the normal reference set $\mathcal{V}_{norm} = \{v_{n_i}\}_{i=1}^N$, where each v_{n_i} denotes the normal-state anchor associated with component v_{p_i} . Each anchor stores the textual standard and,

when available, one or more reference visual features extracted by SAM 3. When multiple normal references are available, they are stored as separate alternatives rather than averaged, preserving intra-class variation during retrieval. Deviations are captured by the defect set $\mathcal{V}_{def} = \{v_{d_j}\}_{j=1}^M$, which enumerates candidate failure modes inferred by GPT-5.

The relational logic is governed by the edge set $\mathcal{E}$, which is partitioned into three functional mappings as

$$\mathcal{E} = \mathcal{E}_{comp} \cup \mathcal{E}_{state} \cup \mathcal{E}_{diff} \quad (2)$$

First, the *hierarchical decomposition edges* $\mathcal{E}_{comp}$ connect the class node v_{cls} to constituent parts in $\mathcal{V}_{part}$ to establish the physical assembly. Second, the *state association edges* $\mathcal{E}_{state}$ branch from each part node v_{p_i} to establish a dual connection comprising a bijective mapping to the unique normal reference v_{n_i} for standardization, and a one-to-many mapping to relevant defect nodes in $\mathcal{V}_{def}$ for anomaly enumeration. Finally, the *differentiation edges* $\mathcal{E}_{diff} \subset \mathcal{V}_{def} \times \mathcal{V}_{def}$ explicitly connect visually confusing defect patterns. These edges are dynamically generated by GPT-5 during the offline knowledge expansion phase to encode contrastive rules (e.g., differentiating a “scratch” from a “crack” via spatial depth cues), thereby enforcing fine-grained discriminative boundaries and mitigating ambiguity during online verification.

3.3 Dual-Stream Verifiable Inference

To reconcile the conflicting requirements of global structural consistency and fine-grained defect sensitivity, we propose a dual-stream architecture where two parallel pathways maximize inspection throughput.

Global & Logic Stream The Global & Logic Stream focuses on detecting macroscopic structural anomalies and logical inconsistencies by performing a preliminary audit of the entire object. This process uses a SAM 3-based fact extractor to transform industrial specifications into partially checkable visual primitives, which effectively grounds the inspection in measurable physical evidence. The global verification is formulated as a structured consistency check defined by

$$\begin{aligned} \mathcal{F}_{fact} &= \Psi_{SAM}(I, \mathcal{R}_{logic}) \\ v_{struct} &= \mathcal{V}_{audit}(I, \mathcal{F}_{fact}, \mathcal{K}_{glob} \mid \mathcal{P}_{chain}) \end{aligned} \quad (3)$$

where I represents the query image and $\mathcal{R}_{logic}$ denotes the set of categorical rules used to guide fact extraction. $\mathcal{R}_{logic}$ is instantiated from PVLA part rules, where predicates are paired with fixed SAM text phrases. Ψ_{SAM} masks the named parts, and the resulting visual rules produce $\mathcal{F}_{fact}$, including component counts, spatial arrangement, color order, and coarse geometry. The term $\mathcal{K}_{glob}$ denotes the multimodal standards, including a background-free global reference image and textual specifications; in zero-shot settings, it reduces to textual specifications only. The structural verdict v_{struct} is an intermediate report rather than the

final diagnosis. It is derived by a multimodal auditor $\mathcal{V}_{\text{audit}}$ operating under a sequential prompting policy $\mathcal{P}_{\text{chain}}$, which first asks the model to identify visual facts from the target image I independently of the standard. These findings are then cross-validated against the logic report $\mathcal{F}_{\text{fact}}$ and the multimodal global knowledge $\mathcal{K}_{\text{glob}}$ to produce the intermediate structural report. By transforming abstract checks into a text-grounded fact-matching task, this stream provides a more grounded and partially checkable basis for the subsequent fine-grained analysis.

Fine-Grained & Actions Stream Exhaustive high-resolution inspection for minute defects is computationally prohibitive. Thus, we formulate fine-grained inspection as a budgeted Markov Decision Process (MDP), where MCTS is used as a budgeted proposal mechanism for selecting evidence crops [8, 15, 19, 34, 44].

The state s_t corresponds to an active image crop parameterized by its bounding box. The action space $\mathcal{A}$ includes discrete translations, scaling, and an inspect region action triggered by global proposals. To guide the search under a fixed budget, we use an anomaly heatmap $\mathcal{H}$ from a lightweight localizer as a heuristic prior. For each state, we compute a heatmap-derived search score S_{search} , which summarizes the local saliency of the crop and its alignment with global proposals. Rather than serving as a new search algorithm, MCTS is used here to maintain exploration over neighboring views and retain spatially diverse evidence crops under the same budget. It performs selection, expansion/simulation, and back-propagation; after exhausting the simulation budget, we apply NMS to obtain final proposals:

1. Selection: Starting from the root, the policy selects child nodes maximizing a modified Upper Confidence Bound (UCB) that balances exploitation, exploration, and the heuristic prior:

$$a_t = \arg \max_{a \in \mathcal{A}} \left(\frac{V(s_{\text{next}})}{N(s_{\text{next}})} + C_p \sqrt{\frac{2 \ln N(s_t)}{N(s_{\text{next}})}} + \lambda S_{\text{search}}(s_{\text{next}}) \right) \quad (4)$$

where $V(\cdot)$ and $N(\cdot)$ are the accumulated value and visit count, C_p is the exploration constant, and $\lambda = 0.5$ weights the heatmap-derived search prior S_{search} to bias the search towards visually deviant areas.

2. Expansion & Simulation: Valid actions expand non-terminal nodes. The new leaf node s_{leaf} is evaluated via a thresholded reward derived from the heatmap:

$$R(s_{\text{leaf}}) = \begin{cases} S_{\text{search}}(s_{\text{leaf}}) & \text{if } S_{\text{search}}(s_{\text{leaf}}) \geq \tau_{\text{sim}} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

where $\tau_{\text{sim}} = 0.05$ is a permissive low-confidence gate for pruning uninformative regions rather than a final anomaly threshold; it is fixed across all experiments.

3. Backpropagation: The scalar reward $R(s_{\text{leaf}})$ is propagated upward through the traversed path to update the node statistics. For each node s in the path, the values are updated as $V(s) \leftarrow V(s) + R(s_{\text{leaf}})$ and $N(s) \leftarrow N(s) + 1$.

4. NMS-based Proposal Filtering: After the budget is exhausted, we apply NMS over visited node coordinates (IoU threshold 0.1) using the MCTS posterior score to rank candidates. The top- K spatially distinct crops form the final evidence set $\mathcal{I}_{crops}$. Compared with directly selecting the maximum or Top- K heatmap peaks, this procedure favors spatially distinct views and yields modest but consistent gains under the same crop budget.

Reference-Guided In-Context Verification In the final stage, the outputs from the parallel streams converge to produce a verified decision. The VLM acts as the final verifier by combining the structural report v_{struct} and the fine-grained evidence crops $\mathcal{I}_{crops}$ with the structured context retrieved from the **Part-Aware Visual-Logical Atlas (PVLA)**. Unlike generic retrieval baselines that rely on flat similarity searches, our mechanism exploits the **hierarchical graph structure** of the PVLA to retrieve component-scoped standards and reference evidence.

The final verdict y is generated by maximizing the consistency between the observed evidence and the retrieved standards via

$$y^* = \arg \max_{y \in \mathcal{Y}} P(y \mid v_{struct}, \mathcal{I}_{crops}, \mathcal{Z}_{graph}, \mathcal{T}_{graph}) \quad (6)$$

In MMAD-QA, $\mathcal{Y}$ is the closed set of answer options for the current question. The verifier is prompted to output the option in the required format, e.g., “The correct answer is (X)”, and the parser extracts X as the final prediction. Thus, the argmax denotes closed-set VLM selection rather than an additional learned optimization module. Here, $P(\cdot)$ formalizes the conditional log-likelihood of the Large Multimodal Model generating the correct diagnostic class sequence y , parameterized by the fused prompt. Crucially, the context retrieval constitutes a topological traversal rather than a simple flat vector search. The system navigates from the root object node down to specific component nodes (e.g., *Cable* $\rightarrow$ *Sheath*) to fetch the precise executable context $\mathcal{T}_{graph}$ and canonical visual normals $\mathcal{Z}_{graph}$ relevant to the current view.

In few-shot settings, this hierarchical alignment allows the model to perform localized comparison against spatially corresponding normal reference samples. Conversely, in zero-shot settings where $\mathcal{Z}_{graph} = \emptyset$, the verification relies on the structured text logic encoded in the leaf nodes. These nodes encode discriminative attributes including the *Visual Signature* and *Contrast Logic* to differentiate between defect types. This allows the VLM to ground evidence crops $\mathcal{I}_{crops}$ into semantic definitions without requiring visual references, while keeping the final output auditable through an evidence chain that points to the visual evidence and failure mode supporting the decision.

4 Experiments

4.1 Setup

We evaluate GLLS on the **MMAD-QA** benchmark [23], specifically selecting the MVTec-AD [6] and VisA [46] subsets. To verify cross-domain robustness, we also

include MPDD [21], DTD [1], and DAGM [37] datasets. Our evaluation focuses on five core metrics relevant to industrial inspection: Anomaly Discrimination, Defect Analysis, Classification, Description, and Localization. Since MMAD-QA is formulated as closed-set visual question answering, many tasks require choosing among fixed answer options rather than producing calibrated anomaly scores. We therefore report task accuracy for the five QA metrics. For cross-dataset binary anomaly-detection experiments, we follow the binary classification accuracy protocol used by the compared methods. Threshold-independent metrics such as AUROC and PRAUC are not uniformly defined for all closed-set QA tasks and are therefore not used as primary metrics. For matched ablations, we fix the VLM backbone, support policy, evidence budget, and verifier so that the effect of PVLA retrieval, dual-stream verification, and local search can be isolated.

Implementation Details We implement GLLS using the PyTorch framework on 8×RTX4090 (48GB) GPUs. For the VLM backbones, we utilize the official instruction-tuned checkpoints of Qwen2.5-VL (7B) and Qwen3-VL (4B, 8B) without any parameter updates. The **Global & Logic Stream** integrates the SAM3 model to execute deterministic segmentation checks. To support the Fine-Grained Action Stream, we employ AdaptCLIP [13] in zero-shot settings and ABound [12] in one-shot settings as the lightweight localizer to generate the anomaly heatmap $\mathcal{H}$ used for heuristic guidance. In the Fine-Grained & Actions Stream, the MCTS is configured with a simulation count of $N = 50$ per decision step and an exploration constant $C_p = 2.0$, with the maximum evidence budget fixed at $K_{max} = 3$ crops based on our ablation analysis. We employ standard chat templates for all VLM interactions. All external modules are frozen and used only as evidence providers or heuristic priors. GPT-5 is used only during offline PVLA construction, SAM 3 provides text-conditioned segmentation for visual facts, and AdaptCLIP/ABound provide heatmap priors for local search. The final diagnostic answer is produced by the VLM verifier from the retrieved PVLA context, structural report, and evidence crops.

4.2 Main Results

Table 1 compares our full system against a diverse set of representative general-purpose Multimodal Large Language Models, including the Qwen-VL, InternVL and MiniCPM-V [2, 3, 35, 38, 40]. GLLS generally improves over the corresponding baselines across both MVTec-AD and VisA subsets. Notably, the gains are not limited to a single stage: the full pipeline improves discrimination reliability while also strengthening defect type classification and region localization. This suggests that the Global & Logic Stream and the Fine-Grained & Actions Stream provide complementary evidence for adapting general VLMs to industrial inspection. Table 1 evaluates general-purpose VLM baselines under the same MMAD-QA protocol, while the ablation tables provide matched comparisons under the same backbone and evidence budget. While we do not explicitly compare against optimization-based methods that have not been benchmarked on MMAD-QA, GLLS offers a competitive and training-free alternative suitable

Table 1: Evaluation on the **MMAD benchmark** [23] across MVTec and VisA domains. Results reflect the accuracy on specific MMAD-QA tasks. Gains and losses relative to the base model are indicated as superscripts (**green** for increase, **red** for decrease). The best results are highlighted in bold.

Model	Param	MVTec						VisA					
		Disc.	Ana.	Cls.	Desc.	Loc.	Avg.	Disc.	Ana.	Cls.	Desc.	Loc.	Avg.
GPT-5-nano	/	73.5	91.8	63.8	75.7	73.9	75.7	70.1	78.0	55.2	72.5	68.8	68.9
GPT-5-mini	/	70.6	93.5	77.5	84.0	82.7	81.7	70.5	83.6	73.3	80.4	78.9	77.3
Gemini-2.5	/	71.8	90.3	72.4	78.7	80.6	78.8	67.1	80.3	68.2	78.7	72.7	73.4
Gemini-3	/	85.7	95.7	86.1	89.9	88.9	89.3	76.4	86.6	80.1	86.3	81.9	82.3
Qwen2.5-VL	3B	68.3	91.5	58.5	76.8	72.5	73.5	50.5	72.0	49.1	74.2	60.8	61.3
InternVL-3.5	4B	74.4	91.5	57.9	78.2	67.7	73.9	73.8	83.7	49.4	71.6	68.8	69.5
MiniCPM-V4	4B	79.4	94.5	68.0	87.2	86.2	83.1	68.9	77.7	51.0	75.3	77.1	70.0
Qwen3-VL	4B	70.5	90.5	64.5	57.5	79.0	72.4	75.0	78.5	55.4	74.6	72.4	71.2
Qwen2.5-VL	7B	75.8	92.3	67.3	82.2	77.2	79.0	63.7	75.6	51.9	71.2	72.7	67.0
InternVL-3.5	8B	66.8	93.3	64.5	81.6	75.9	76.4	61.4	80.7	52.8	76.1	74.6	69.1
Qwen3-VL	8B	83.1	93.5	74.2	84.8	86.6	84.4	69.1	81.8	61.4	77.5	81.2	74.2
Qwen3-VL-Thinking	8B	81.7	93.7	69.7	83.5	81.7	82.1	66.4	78.8	59.3	79.6	76.2	72.1
MiniCPM-V4.5	9B	77.8	93.5	72.5	89.4	80.3	82.7	61.0	81.9	61.1	79.0	70.1	70.6
GLLS (Qwen3-VL, 0-shot)	4B	82.6 ^{+12.1}	91.3 ^{+0.8}	79.8 ^{+15.3}	77.3 ^{+19.8}	77.1 ^{-1.9}	81.6 ^{+9.2}	75.4 ^{+0.4}	78.2 ^{-0.8}	71.5 ^{+16.1}	78.7 ^{+4.1}	78.5 ^{+6.1}	76.5 ^{+5.3}
GLLS (Qwen2.5-VL, 0-shot)	7B	84.0 ^{+8.2}	92.4 ^{+0.1}	73.3 ^{+6.0}	83.6 ^{+1.4}	77.4 ^{+0.2}	82.1 ^{+3.1}	72.6 ^{+8.9}	75.7 ^{+0.1}	62.2 ^{+10.3}	71.0 ^{-0.2}	71.6 ^{+1.1}	70.6 ^{+3.6}
GLLS (Qwen3-VL, 0-shot)	8B	88.6 ^{+5.5}	93.9 ^{+0.4}	81.9 ^{+7.7}	85.0 ^{+0.2}	86.2 ^{-0.4}	87.1 ^{+2.7}	74.3 ^{+5.2}	80.9 ^{-0.9}	65.2 ^{+3.8}	75.5 ^{-2.0}	79.8 ^{+1.4}	75.1 ^{+0.9}
GLLS (Qwen3-VL, 1-shot)	4B	96.4 ^{+25.9}	92.5 ^{+2.0}	84.3 ^{+19.8}	81.9 ^{+24.4}	83.2 ^{+4.2}	87.7 ^{+15.3}	87.4 ^{+12.4}	81.7 ^{+3.2}	73.8 ^{+18.4}	75.6 ^{+1.0}	81.6 ^{+9.2}	80.0 ^{+8.8}
GLLS (Qwen2.5-VL, 1-shot)	7B	93.1 ^{+17.3}	92.7 ^{+0.4}	78.2 ^{+10.9}	83.5 ^{+1.3}	79.8 ^{+2.6}	85.5 ^{+6.5}	82.9 ^{+19.2}	81.6 ^{+6.0}	68.4 ^{+16.5}	76.2 ^{+5.0}	80.7 ^{+8.0}	78.0 ^{+11.0}
GLLS (Qwen3-VL, 1-shot)	8B	96.5 ^{+13.4}	95.5 ^{+2.0}	87.7 ^{+13.5}	86.9 ^{+2.1}	89.3 ^{+2.7}	91.2 ^{+6.8}	88.2 ^{+19.1}	84.7 ^{+2.9}	83.5 ^{+22.1}	80.2 ^{+2.7}	87.0 ^{+5.8}	84.7 ^{+10.5}

for rapid deployment. To further evaluate the robustness of GLLS, we extend the evaluation to MPDD, DTD, and DAGM (Table 2). In the 1-shot setting, GLLS achieves an average accuracy of **91.01%**, outperforming proprietary LMMs and several specialized CLIP-based visual anomaly detectors under this binary accuracy protocol [12, 13, 20, 45]. This indicates that the proposed training-free pipeline leverages visual-logical cues that transfer across datasets.

4.3 Ablation Studies and Component Analysis

We performed a systematic evaluation to isolate the contributions of individual components and to assess the trade-off between performance and computational cost.

Synergistic Contribution of Components Table 3 presents the performance impact of removing key modules. The offline-only strategy, which uses static knowledge retrieval without active perception, often improves the base model but exhibits instability on the VisA dataset due to visual hallucinations. Conversely, the online-only configuration using dual-stream verification lacks the semantic grounding provided by the PVLA, and consequently struggles with defect classification. The full GLLS framework integrates both streams to achieve the best average results in most settings, indicating that rigorous offline standards benefit from active online verification.

Table 4a demonstrates the incremental benefits of the online dual-stream architecture. Adding the *Fine-Grained & Actions Stream* to the offline baseline yields consistent gains by acquiring high-resolution details, and the full GLLS achieves the highest performance. This result suggests that local evidence is more effective when contextualized by global reasoning.

Furthermore, Table 4b isolates the contribution of the offline PVLA. Although PVLA generally outperforms uni-modal baselines, the performance drop observed

Table 2: Cross-dataset generalization performance in industrial anomaly detection excluding VisA. We report binary classification accuracy (%). * denotes threshold optimized using Youden’s J statistic. **GLLS (Ours)** variants are listed in the final section.

	Model	Scale	MPDD	DTD	DAGM	Avg.
CLIP-based	WinClip [20] ($\tau = 0.5$)	-	72.10	25.90	86.90	57.08
	WinClip [20]*	-	41.80	90.00	88.90	75.85
	AnomalyClip [45] ($\tau = 0.5$)	-	56.49	81.71	22.19	54.24
	AnomalyClip [45]*	-	68.23	88.91	95.36	83.67
	AdaptCLIP [13] ($\tau = 0.5$)	-	63.98	84.00	82.17	76.72
	AdaptCLIP [13]*	-	72.16	91.89	94.37	86.14
Proprietary	GPT4o-mini	-	67.90	90.56	90.80	80.81
	GPT4o	-	69.43	90.80	90.17	81.04
	GPT-4.1	-	66.81	90.18	89.93	81.98
Open-source	LLaVA-NeXT	34B	61.57	72.62	23.15	53.70
	InternVL2	76B	66.38	90.58	94.52	80.32
	InternVL3	8B	70.74	85.05	81.33	75.54
	InternVL3	38B	69.43	90.79	89.87	79.71
	Qwen2.5-VL	7B	68.34	85.59	94.83	76.07
Fine-tuned	IAD-R1 [27] (LLaVA-OV)	7B	70.90	96.20	94.80	86.10
	AD-FM [29](Qwen2.5-VL)	7B	72.71	92.64	95.51	84.57
Ours	GLLS (Qwen3-VL, <i>0-shot</i>)	4B	72.34	93.12	96.32	87.26
	GLLS (Qwen2.5-VL, <i>0-shot</i>)	7B	71.56	88.34	95.43	85.11
	GLLS (Qwen3-VL, <i>0-shot</i>)	8B	75.41	93.40	97.15	88.65
	GLLS (Qwen3-VL, <i>1-shot</i>)	4B	78.45	93.56	96.65	89.55
	GLLS (Qwen2.5-VL, <i>1-shot</i>)	7B	73.10	91.45	96.12	86.89
	GLLS (Qwen3-VL, <i>1-shot</i>)	8B	80.23	95.31	97.48	91.01

with Qwen2.5-7B reveals a specific limitation: enforcing strict standards on passive, low-resolution inputs increases false positives. This finding indicates that rigorous offline knowledge depends on active online verification to successfully distinguish defects from noise.

System Efficiency and Search Robustness Table 5 reports the latency for the *Global & Logic* and *Fine-Grained & Actions* streams (denoted as *Global Logic* and *Local Search*, respectively). Results show that the latter operates faster than the former, so the additional search overhead remains small relative to final VLM verification. We further analyze the MCTS components in Table 6: **(a) Search Policy:** MCTS-driven planning improves over sliding window and random baselines by leveraging heatmap heuristics to prioritize informative regions. **(b) Evidence Budget:** The VLM budget K_{max} controls the cost-performance trade-off. We find $K_{max} = 3$ to be a good operating point, as higher values increase latency with diminishing returns. **(c, d) Stability & Efficiency:** The system shows limited sensitivity to these hyperparameters. Varying the simulation count or exploration constant results in marginal fluctuations. As shown in column (c), even with $N = 100$, the search latency is only 0.15s, keeping the planning overhead small.

Table 3: Ablation study of the proposed GLLS method on MVTec and VisA datasets. We compare the Base model, Offline-only (Knowledge Retrieval), Online-only (Two-stream Verification), and the full GLLS method. The checkmark (✓) indicates the inclusion of the corresponding module. Best results for each model are bolded.

Model	Param	Strategy		MVTec						VisA					
		Offline	Online	Disc.	Ana.	Cls.	Desc.	Loc.	Avg.	Disc.	Ana.	Cls.	Desc.	Loc.	Avg.
Qwen3-VL-Instruct	4B	✓		70.5	90.5	64.5	57.5	79.0	72.40	75.0	78.5	55.4	74.6	72.4	71.18
			✓	88.7	92.4	84.1	81.8	81.6	85.72	74.7	81.5	72.3	77.6	76.9	76.60
		✓	✓	96.4	92.5	84.3	81.9	83.2	87.66	87.4	81.7	73.8	75.6	81.6	80.02
Qwen2.5-VL-Instruct	7B	✓		75.8	92.3	67.3	82.2	77.2	78.96	63.7	75.6	51.9	71.2	72.7	67.02
			✓	80.2	86.8	74.9	74.4	68.9	77.04	64.0	68.0	57.9	64.0	60.9	62.96
		✓	✓	93.1	92.7	78.2	83.5	79.8	85.46	82.9	79.6	68.4	76.2	80.7	77.56
Qwen3-VL-Instruct	8B	✓		83.1	93.5	74.2	84.8	86.6	84.44	69.1	81.8	61.4	77.5	81.2	74.20
			✓	89.7	94.4	87.5	85.8	87.5	88.98	77.7	81.5	78.9	79.5	81.7	79.86
		✓	✓	96.5	95.5	87.7	86.9	89.3	91.18	88.2	84.7	83.5	80.2	87.0	84.72

Table 4: Ablation Studies on Core Architectural Components. (a) Step-wise ablation demonstrating the contribution of the Online Dual-Stream architecture. (b) Ablation on the Part-Aware Visual-Logical Atlas (PVLA) isolating the contribution of the Offline Construction strategy.

(a) Impact of Online Dual-Stream					(b) Impact of Offline PVLA				
Model	Config.	MVTec	VisA	Avg.	Model	Method	MVTec	VisA	Avg.
Qwen3 4B	Base (Offline)	85.7	76.6	81.2	Qwen3 4B	Base	72.4	71.2	71.8
	+ Fine-Grained	86.8	78.5	82.7		Text-Guided	80.5	73.2	76.9
	GLLS	87.7	80.0	83.8		Image-Guided	82.1	74.1	78.1
Qwen2.5 7B	Base (Offline)	77.0	63.0	70.0	Qwen2.5 7B	PVLA	85.7	76.6	81.2
	+ Fine-Grained	82.3	73.1	77.7		Base	79.0	67.0	73.0
	GLLS	85.5	77.6	81.5		Text-Guided	77.8	65.5	71.7
Qwen3 8B	Base (Offline)	89.0	79.9	84.4	Qwen3 8B	Image-Guided	76.5	64.1	70.3
	+ Fine-Grained	89.9	82.6	86.2		PVLA	77.0	63.0	70.0
	GLLS	91.2	84.7	88.0		Base	84.4	74.2	79.3
						Text-Guided	86.5	76.4	81.5
						Image-Guided	87.2	78.1	82.7
						PVLA	89.0	79.9	84.4

4.4 Qualitative Visualization

To demonstrate the interpretability of our framework, Figure 3 visualizes a zipper inspection case involving a Split Teeth defect. In the Global & Logic Stream, SAM 3 segments key components such as the zipper teeth, and the resulting visual primitives undergo deterministic PVLA checks. The logical report is combined with global multimodal knowledge and fed into the VLM to derive a preliminary conclusion on structural integrity. Concurrently, the Fine-Grained & Actions Stream uses MCTS to capture high-resolution evidence crops revealing the separation. Final verification fuses global structural assessment with local visual details to produce an evidence-supported verdict traceable to visual observations.

Table 5: Efficiency Analysis. Latency breakdown. Total latency $T_{total} \approx \max(T_{glob}, T_{search}) + T_{verif}$ shows that the fine-grained search overhead is largely masked by the parallel logic checks. **G&L**: Global & Logic; **FG&A**: Fine-Grained & Actions.

Model	Method	Dual Streams		Reasoning	Total
		<i>G&L</i>	<i>FG&A</i>	<i>Verification</i>	
Qwen3 4B	Base	-	-	0.74s	0.74s
	GLLS (0-shot)	0.33s	0.04s	0.77s	1.10s
	GLLS (1-shot)	0.31s	0.07s	0.90s	1.21s
Qwen3 8B	Base	-	-	0.77s	0.77s
	GLLS (0-shot)	0.31s	0.04s	1.00s	1.31s
	GLLS (1-shot)	0.33s	0.08s	1.35s	1.68s
Qwen2.5 7B	Base	-	-	0.85s	0.85s
	GLLS (0-shot)	0.38s	0.03s	1.00s	1.38s
	GLLS (1-shot)	0.32s	0.08s	1.52s	1.84s

Table 6: Ablation Study on Fine-Grained & Actions Stream. Default setting ($K_{max} = 3, N = 50, C_p = 2.0$) is highlighted in gray. We report *Total Time* for (b) and specific *Search Time* for (c).

(a) Search Policy				(b) Evidence Budget (K_{max})			
Strategy	MVTec	VisA	IoU	Budget	MVTec	VisA	Total Time
Center Crop	84.2	75.1	65.2	1 Crop	88.5	81.2	1.15s
Random	86.5	78.3	69.8	3 Crops	91.2	84.7	1.68s
Sliding Window	89.4	82.4	74.5	5 Crops	90.3	83.1	2.24s
GLLS (MCTS)	91.2	84.7	78.5	7 Crops	90.1	82.0	2.89s
(c) Simulation Count (N)				(d) Exploration Const. (C_p)			
Count	MVTec	VisA	Search Time	Value	MVTec	VisA	IoU
$N = 10$	90.4	84.1	0.02s	$C_p = 0.5$	90.5	84.0	77.6
$N = 30$	90.8	84.5	0.05s	$C_p = 1.0$	90.9	84.3	78.1
$N = 50$	91.2	84.7	0.08s	$C_p = 2.0$	91.2	84.7	78.5
$N = 100$	91.1	84.6	0.15s	$C_p = 4.0$	90.6	84.2	77.9

5 Conclusion

We presented GLLS, a training-free dual-stream framework for verifiable industrial anomaly detection. The core idea is verification-first grounding, in which the system casts LMMs as evidence-guided auditors by providing part-level standards and making intermediate signals partially checkable through frozen external tools. By combining **Global & Logic** for structural compliance with **Fine-Grained & Actions** for detailed visual evidence, GLLS achieves strong performance on the MMAD-QA benchmark and shows cross-dataset generalization across several industrial domains including MPDD, DTD, and DAGM. Future work will focus on learning the search policy with verifiable rewards and expanding PVLA with richer contrastive edges to further improve discrimination accuracy.




Fig. 3: Qualitative Visualization of the GLLS Inference Process. Using a zipper inspection as an example, the framework illustrates the synthesis of global structural checks and local high-resolution evidence for verifiable diagnosis.

Acknowledgement.

This work was supported by the National Natural Science Foundation of China (No. 62476124, 62302093), the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (JYB2025XDXM902), the Natural Science Foundation of Jiangsu Province (No. BK20242015, BK20230833), the Gusu Innovation and Entrepreneur Leading Talents (No. ZXL2025322), and Nanjing University-China Mobile Communications Group Co.,Ltd. Joint Institute.

References

1. Aota, T., Tong, L.T.T., Okatani, T.: Zero-shot versus many-shot: Unsupervised texture anomaly detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5564–5572 (2023)

2. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., et al.: Qwen3-vl technical report. arXiv preprint arXiv:2511.21631 (2025)
3. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., et al.: Qwen2.5-vl technical report. arXiv preprint arXiv:2502.13923 (2025)
4. Bajcsy, R., Aloimonos, Y., Tsotsos, J.K.: Revisiting active perception. *Autonomous Robots* **42**(2), 177–196 (2018)
5. Baldassini, F.B., Shukor, M., Cord, M., Soulier, L., Piwowarski, B.: What makes multimodal in-context learning work? In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1539–1550 (2024)
6. Bergmann, P., Fauser, M., Sattlegger, D., Steger, C.: Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9592–9600 (2019)
7. Bigverdi, M., Luo, Z., Hsieh, C.Y., Shen, E., Chen, D., Shapiro, L.G., Krishna, R.: Perception tokens enhance visual reasoning in multimodal language models. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 3836–3845 (2025)
8. Browne, C.B., Powley, E., Whitehouse, D., Lucas, S.M., Cowling, P.I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., Colton, S.: A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games* **4**(1), 1–43 (2012)
9. Carion, N., Gustafson, L., Hu, Y.T., Debnath, S., Hu, R., Suris, D., Ryali, C., Alwala, K.V., Khedr, H., Huang, A., et al.: Sam 3: Segment anything with concepts. arXiv preprint arXiv:2511.16719 (2025)
10. Chao, Y., Liu, J., Tang, J., Wu, G.: Anomalyr1: A grpo-based end-to-end mllm for industrial anomaly detection. arXiv preprint arXiv:2504.11914 (2025)
11. Chen, Z., Chen, H., Imani, M., Imani, F.: Can multimodal large language models be guided to improve industrial anomaly detection? In: *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. vol. 89213, p. V02BT02A051. American Society of Mechanical Engineers (2025)
12. Deng, R., Hu, Y., Zhang, X., Wang, Z., Liu, X., Dai, W.Z., Shan, C., Zhao, F.: Abound: Adversarial boundary-driven few-shot learning for multi-class anomaly detection. arXiv preprint arXiv:2511.22436 (2025)
13. Gao, B.B., Zhou, Y., Yan, J., Cai, Y., Zhang, W., Wang, M., Liu, J., Liu, Y., Wang, L., Wang, C.: Adaptclip: Adapting clip for universal visual anomaly detection. arXiv preprint arXiv:2505.09926 (2025)
14. Guan, W., Lan, J., Cao, J., Tan, H., Zhu, H., Wang, W.: Emit: Enhancing mllms for industrial anomaly detection via difficulty-aware grpo. arXiv preprint arXiv:2507.21619 (2025)
15. Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q., Xu, R., Zhang, R., Ma, S., Bi, X., et al.: Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948 (2025)
16. Hochstein, S., Ahissar, M.: View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron* **36**(5), 791–804 (2002)
17. Huang, B., Mitra, C., Arbelle, A., Karlinsky, L., Darrell, T., Herzig, R.: Multimodal task vectors enable many-shot multimodal in-context learning. *Advances in Neural Information Processing Systems* **37**, 22124–22153 (2024)

18. Hurst, A., Lerer, A., Goucher, A.P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al.: Gpt-4o system card. arXiv preprint arXiv:2410.21276 (2024)
19. Jaech, A., Kalai, A., Lerer, A., Richardson, A., El-Kishky, A., Low, A., Helyar, A., Madry, A., Beutel, A., Carney, A., et al.: Openai o1 system card. arXiv preprint arXiv:2412.16720 (2024)
20. Jeong, J., Zou, Y., Kim, T., Zhang, D., Ravichandran, A., Dabeer, O.: Winclip: Zero-/few-shot anomaly classification and segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 19606–19616 (2023)
21. Jezek, S., Jonak, M., Burget, R., Dvorak, P., Skotak, M.: Deep learning-based defect detection of metal parts: evaluating current methods in complex conditions. In: 2021 13th International congress on ultra modern telecommunications and control systems and workshops (ICUMT). pp. 66–71. IEEE (2021)
22. Ji, D., Hu, B., Zhou, Y.: Autoiad: Manager-driven multi-agent collaboration for automated industrial anomaly detection. arXiv preprint arXiv:2508.05503 (2025)
23. Jiang, X., Li, J., Deng, H., Liu, Y., Gao, B.B., Zhou, Y., Li, J., Wang, C., Zheng, F.: Mmad: A comprehensive benchmark for multimodal large language models in industrial anomaly detection. arXiv preprint arXiv:2410.09453 (2024)
24. Jin, E., Feng, Q., Mou, Y., Lakemeyer, G., Decker, S., Simons, O., Stegmaier, J.: Logicad: Explainable anomaly detection via vlm-based text feature extraction. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 4129–4137 (2025)
25. Li, W., Fan, H., Wong, Y., Yang, Y., Kankanhalli, M.S.: Improving context understanding in multimodal large language models via multimodal composition learning. In: ICML. vol. 3, p. 7 (2024)
26. Li, W., Chu, G., Chen, J., Xie, G.S., Shan, C., Zhao, F.: Lad-reasoner: Tiny multimodal models are good reasoners for logical anomaly detection. arXiv preprint arXiv:2504.12749 (2025)
27. Li, Y., Cao, Y., Liu, C., Xiong, Y., Dong, X., Huang, C.: Iad-r1: Reinforcing consistent reasoning in industrial anomaly detection. arXiv preprint arXiv:2508.09178 (2025)
28. Li, Y., Yuan, S., Wang, H., Li, Q., Liu, M., Xu, C., Shi, G., Zuo, W.: Triad: Empowering lmm-based anomaly detection with expert-guided region-of-interest tokenizer and manufacturing process. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21917–21926 (2025)
29. Liao, J., Su, Y., Tu, R.C., Jin, Z., Sun, W., Li, Y., Tao, D., Xu, X., Yang, X.: Ad-fm: Multimodal llms for anomaly detection via multi-stage reasoning and fine-grained reward optimization. arXiv preprint arXiv:2508.04175 (2025)
30. Mecklenburg, N., Lin, Y., Li, X., Holstein, D., Nunes, L., Malvar, S., Silva, B., Chandra, R., Aski, V., Yannam, P.K.R., et al.: Injecting new knowledge into large language models via supervised fine-tuning. arXiv preprint arXiv:2404.00213 (2024)
31. Miao, J., Du, P., Liu, Y., Wang, Y., Wang, Y.: Agentiad: Tool-augmented single-agent for industrial anomaly detection. arXiv preprint arXiv:2512.13671 (2025)
32. OpenAI: Introducing gpt-5. <https://openai.com/zh-Hans-CN/index/introducing-gpt-5/> (August 7 2025), accessed: 2026-01-27
33. Song, Z., Yan, B., Liu, Y., Fang, M., Li, M., Yan, R., Chen, X.: Injecting domain-specific knowledge into large language models: A comprehensive survey (2025), <https://arxiv.org/abs/2502.10708>

34. Wang, H., Su, A., Ren, W., Lin, F., Chen, W.: Pixel reasoner: Incentivizing pixel-space reasoning with curiosity-driven reinforcement learning. arXiv preprint arXiv:2505.15966 (2025)
35. Wang, W., Gao, Z., Gu, L., Pu, H., Cui, L., Wei, X., Liu, Z., Jing, L., Ye, S., Shao, J., et al.: Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. arXiv preprint arXiv:2508.18265 (2025)
36. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* **35**, 24824–24837 (2022)
37. Wieler, M., Hahn, T.: Weakly supervised learning for industrial optical inspection. In: DAGM symposium in. vol. 6, p. 11 (2007)
38. Yao, Y., Yu, T., Zhang, A., Wang, C., Cui, J., Zhu, H., Cai, T., Li, H., Zhao, W., He, Z., et al.: Minicpm-v: A gpt-4v level mllm on your phone. arXiv preprint arXiv:2408.01800 (2024)
39. Yu, S., Tang, C., Xu, B., Cui, J., Ran, J., Yan, Y., Liu, Z., Wang, S., Han, X., Liu, Z., et al.: Visrag: Vision-based retrieval-augmented generation on multi-modality documents. arXiv preprint arXiv:2410.10594 (2024)
40. Yu, T., Wang, Z., Wang, C., Huang, F., Ma, W., He, Z., Cai, T., Chen, W., Huang, Y., Zhao, R., et al.: Minicpm-v 4.5: Cooking efficient mllms via architecture, data, and training recipe. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11704–11715 (2026)
41. Zeng, P., Pang, F., Wang, Z., Yang, A.: Lr-iad: Mask-free industrial anomaly detection with logical reasoning. arXiv preprint arXiv:2504.19524 (2025)
42. Zhang, Z., Ruan, J., Gao, X., Liu, T., Fu, Y.: Eiad: Explainable industrial anomaly detection via multi-modal large language models. In: *2025 IEEE International Conference on Multimedia and Expo (ICME)*. pp. 1–6. IEEE (2025)
43. Zhao, S., Lin, Y., Han, L., Zhao, Y., Wei, Y.: Omniad: Detect and understand industrial anomaly via multimodal reasoning. arXiv preprint arXiv:2505.22039 (2025)
44. Zheng, Z., Yang, M., Hong, J., Zhao, C., Xu, G., Yang, L., Shen, C., Yu, X.: Deepeyes: Incentivizing "thinking with images" via reinforcement learning. arXiv preprint arXiv:2505.14362 (2025)
45. Zhou, Q., Pang, G., Tian, Y., He, S., Chen, J.: Anomalyclip: Object-agnostic prompt learning for zero-shot anomaly detection. In: *International Conference on Learning Representations*. vol. 2024, pp. 49705–49737 (2024)
46. Zou, Y., Jeong, J., Pemula, L., Zhang, D., Dabeer, O.: Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In: *European conference on computer vision*. pp. 392–408. Springer (2022)